arXiv:1701.04475v1 [math.CO] 16 Jan 2017

THE MULTI-SLICE RANK METHOD AND POLYNOMIAL BOUNDS FOR ORTHOGONAL SYSTEMS IN $\mathbb{F}_q^n$

ERIC NASLUND

ABSTRACT. Let $\mathcal{S} \subset \mathbb{F}_q^n$ be the subset of self-orthogonal vectors in $\mathbb{F}_q^n$, which has size $|\mathcal{S}| = q^{n-1} + O\left(\sqrt{q^n}\right)$. In this paper we use the recently developed slice rank method [11] to show that any set $E \subset \mathcal{S}$ of size

$$|E| > (n+1)^{(q-1)(k-1)}$$

contains a k-tuple of distinct mutually orthogonal vectors, where $n \geq \binom{k}{2}$ and $q = p^r$ with $p \geq k$. The key innovation is a more general version of the slice rank of a tensor, which we call the *multi-slice rank*. We use a combinatorial argument to create an appropriate indicator tensor for the orthogonality of k-vectors, however for $k \geq 4$ this tensor will have large slice rank, and so to obtain our results we need to use the less restrictive multi-slice rank. Additionally, we use this method to generalize a recent result of Ge and Shangguan [7], and prove that any set $A \subset \mathbb{F}_q^n$ of size

$$|A| > \binom{n + (q-1)(k-1)}{(q-1)(k-1)}$$

contains a *k-right-corner*, that is distinct vectors $x_1, \ldots, x_k, x_{k+1}$ where $x_1 - x_{k+1}, \ldots, x_k - x_{k+1}$ are mutually orthogonal, as long as $n \geq \binom{k+1}{2}$ and $q = p^r$ with $p > k$.

1. Introduction

Problem 1. *What is the size of the largest subset $E \subset \mathbb{F}_q^n$ which contains no distinct k-tuple of mutually orthogonal vectors?*

In 2008 Vinh [13] generalized a result of Iosevich and Senger [8], proving that for $n \geq 2k-1$, if $E \subset \mathbb{F}_q^n$ satisfies

$$|E| \gg q^{\frac{n}{2}+k-1},$$

then E contains

$$(1 + o(1)) \frac{|E|^k}{k!} q^{-\binom{k}{2}}$$

mutually orthogonal k-tuples in E. Here $q = p^r$ is a prime power, and $\mathbb{F}_q^n$ denotes the vector space of dimension n over the finite field with q elements. In Iosevich and Senger's [8, Lemma 3.2] paper, they constructed a set $E \subset \mathbb{F}_q^n$ of size

$$|E| \geq c q^{\frac{n}{2}+1}$$

for some $c > 0$ such that no two of it's vectors are orthogonal. In light of Vinh's bound, when $k = 2$ this construction is tight up to a constant, and so it seems that problem 1 is resolved in this case. However, critical to this lower bound construction is the fact that none of the vectors in this set are self-orthogonal, that is no $x \in E$ satisfies $\sum_{i=1}^n x_i^2 = 0$. Let $\mathcal{S}$ denote the set of all self-orthogonal vectors in $\mathbb{F}_q^n$, also referred to as the set of *isotropic* vectors (see [1]). We note that

$$|\mathcal{S}| = q^{n-1} + \frac{1}{q} \sum_{\substack{\chi \in \hat{\mathbb{F}}_q \\ \chi \neq \chi_0}} \left(\sum_{x \in \mathbb{F}_q} \chi(x^2) \right)^n$$

Date: January 12 2017.

where χ ranges over the characters of $\mathbb{F}_q$, and so

$$|\mathcal{S}| = q^{n-1} + O\left(q^{\frac{n}{2}}\right)$$

since the Gauss sum $\sum_{x\in\mathbb{F}_q}\chi(x^2)$ has magnitude $q^{\frac{1}{2}}$. Our main theorem is:

Theorem 1. *Let $\mathcal{S}\subset\mathbb{F}_q^n$ be the set of all $x\in\mathbb{F}_q^n$ satisfying $\sum_{i=1}^n x_i^2 = 0$. Let $k>1$ be given, and suppose that $n\geq\binom{k}{2}$ and $q=p^r$ satisfies $p\geq k$. If $E\subset\mathcal{S}$ has size*

$$|E| > \binom{n+(q-1)(k-1)}{(q-1)(k-1)},$$

then E contains a k-tuple of k distinct mutually orthogonal vectors.

The right hand side above is a polynomial in the dimension and so the result holds for very sparse subsets of $\mathcal{S}$. In particular, the right hand side is bounded above by

$$(n+1)^{(q-1)(k-1)},$$

which is at most

$$O_{\epsilon,k,q}(1+\epsilon)^n$$

for any $\epsilon>0$, where the subscript denotes that the implicit constant could depend on ϵ, k and q. To prove theorem 1, we use the *slice rank method* which was introduced by Tao [11] following the breakthrough of Croot, Lev and Pach [3]. The slice rank method for 3-tensors is a symmetrized version of Croot, Lev and Pach's argument, and in his blog post, Tao used it to reprove Ellenberg and Gijswijt's bounds for arithmetic progression-free subsets of $\mathbb{F}_p^n$ [4]. This method has seen applications to a variety of problems including the Erdő-Szemerédi sunflower problem [10, 6], and the group theoretic approach to fast matrix multiplication [2], and we refer the reader to [2, Section 4] for an in depth discussion of the slice rank as well as its connection to geometric invariant theory. The slice rank method allows us to prove theorem 1 for $k=2$ and $k=3$, however to handle $k\geq 4$ we need additional flexibility, and we introduce a generalization of the slice rank which we call the *multi-slice rank*. Using a combinatorial argument coming from the group action of S_k on the set of k-tuples of vectors in $\mathbb{F}_q^n$, we create a tensor I_k satisfying

$$I_k(x_1,\dots,x_k) = \begin{cases} 1 & x_1,\dots,x_k \text{ are distinct and mutually orthogonal} \\ (-1)^k(k-1)! & \text{if } x_1=\cdots=x_k \\ 0 & \text{otherwise} \end{cases},$$

so that I_k is diagonal when restricted to $E\times\cdots\times E$, where E is any set that does not contain a k-tuple of distinct mutually orthogonal vectors. For $k\geq 4$, I_k will have very large slice rank, and so the slice rank method fails to obtain non-trivial bounds. However I_k will have polynomially small multi-slice rank, and so the multi-slice rank method allows us to prove theorem 1 for $k\geq 4$.

Recently, Ge and Shangguan [7] used the slice rank method to improve upon a result of Bennett [1], showing that any set $A\subset\mathbb{F}_q^n$ of size

$$|A| > \binom{n+q}{q-1} + 3$$

contains a *right-angle*, that is distinct vectors $x,y,z\in A$ satisfying

$$\langle y-x, z-x\rangle = 0.$$

Using the multi-slice rank method we generalize this result to *k-right-corners*. We say that the vectors $x_1,\dots,x_k,x_{k+1}$ form a *k-dimensional right corner* if the k vectors $x_1-x_{k+1},\dots,x_k-x_{k+1}$ form a mutually orthogonal k-tuple. In other words, $x_1,\dots,x_k$ must meet x_{k+1} at mutually right angles.

Theorem 2. *Let k be given, $q = p^r$ with $p > k$, and $n \geq \binom{k+1}{2}$. If $A \subset \mathbb{F}_p^n$ has size*

$$|A| > \binom{n + (q-1)(k-1)}{(q-1)(k-1)},$$

then A contains a k-dimensional right corner.

The proof of theorem 2 relies heavily on the flexibility of the multi-slice rank for $k \geq 3$. Just as before, when we have 4 or more variables involved, the constructed indicator tensor will have polynomial small multi-slice rank yet exponentially large slice rank. In [9], we provide another application of the multi-slice rank method, and prove exponential upper bounds for the Erdős-Ginzburg-Ziv [5] constant of $(\mathbb{Z}/p\mathbb{Z})^n$.

In subsection 2.1, we review the slice rank, as defined in [11] and [2, Section 4.1], and in subsection 2.2 we introduce the multi-slice rank and prove the critical lemma, which states that the multi-slice rank of a diagonal tensor equals the number of non-zero diagonal entries. In 3.1 we introduce the indicator function $H_k(x_1, \ldots, x_k)$, which satisfies

$$H_k(x_1, \ldots, x_k) = \begin{cases} 1 & \text{if } x_1, \ldots, x_k \text{ are distinct} \\ (-1)^k(k-1)! & \text{if } x_1 = \cdots = x_k \\ 0 & \text{otherwise} \end{cases}.$$

This function allows us to modify our tensor so that it picks up only k-tuples of distinct elements. In subsections 3.2 and 3.3, we prove theorems 1 and 2, respectively, using the multi-slice rank and the indicator H_k.

2. The Slice Rank and Multi-slice rank

2.1. **The Slice Rank.** We begin by recalling the tensor rank and the definition of the rank of a two variable function. Let X, Y be finite sets, and suppose that $\mathbb{F}$ is a field. The rank of a function

$$F : X \times Y \to \mathbb{F}$$

is defined to be the smallest k such that

$$F(x, y) = \sum_{i=1}^{k} f_i(x) g_i(y)$$

for some functions f_i, g_i. For finite sets $X_1, \ldots, X_n$, a function

$$h : X_1 \times \cdots \times X_n \to \mathbb{F},$$

$h \neq 0$, of the form

$$h(x_1, \ldots, x_n) = f_1(x_1) f_2(x_2) \cdots f_n(x_n)$$

is called a *rank 1 function,* and the tensor rank of

$$F : X_1 \times \cdots \times X_n \to \mathbb{F}$$

is defined to be the minimal k such that

$$F = \sum_{i=1}^{k} g_i$$

where the g_i are rank 1 functions. Following the breakthrough of Croot, Lev and Pach [3], Tao [11] introduced the notion of the slice rank of a tensor:

Definition 1. Let $X_1, \ldots, X_n$ be finite sets. We say that the function

$$h : X_1 \times \cdots \times X_n \to \mathbb{F},$$

$h \neq 0$, is a *1-slice* if

$$h(x_1, \ldots, x_n) = f(x_i) g\left(x_1, \ldots, x_{i-1}, x_{i+1}, \ldots, x_n\right)$$

for some $1 \leq i \leq n$. The *slice rank*, or *1-slice rank*, of

$$F : X_1 \times \cdots \times X_n \to \mathbb{F}$$

is the smallest k such that

$$F = \sum_{i=1}^{k} g_i$$

where the g_i are 1-slices.

The following lemma was proven by Tao, and used to great effect:

Lemma 1. *Let X be a finite set, and let X^k denote the k-fold Cartesian product of X with itself. Suppose that*

$$F : X^k \to \mathbb{F}$$

is a diagonal tensor, that is

$$F(x_1, \ldots, x_n) = \sum_{a \in A} c_a \delta_a(x_1) \cdots \delta_a(x_n)$$

for some $A \subset X$, $c_a \neq 0$, where

$$\delta_a(x) = \begin{cases} 1 & x = a \\ 0 & otherwise \end{cases}.$$

Then

$$\text{slice-rank}(F) = |A|.$$

Proof. We refer the reader to [11, Lemma 1] or [2, Lemma 4.7]. □

2.2. **The Multi-Slice Rank.** In what follows, we propose a new more general definition of the slice rank. Given variables $x_1, \ldots, x_n$ and a set $S \subset \{1, \ldots, n\}$, $S = \{s_1, \ldots, s_k\}$, we use the notation $\overrightarrow{x_S}$ to denote

$$x_{s_1}, \ldots, x_{s_k},$$

and so for a function g of k variables, we have that

$$g(\overrightarrow{x_S}) = g(x_{s_1}, \ldots, x_{s_k}).$$

Definition 2. Let $X_1, \ldots, X_n$ be finite sets. We say that the function

$$h : X_1 \times \cdots \times X_n \to \mathbb{F},$$

$h \neq 0$, is a *b-slice* for $1 \leq b \leq \frac{n}{2}$ if

$$h(x_1, \ldots, x_n) = f(\overrightarrow{x_S}) g\left(\overrightarrow{x}_{\{1,\ldots,n\} \setminus S}\right)$$

for some f, g and $S \subset \{1, 2, \ldots, n\}$ of size $|S| = b$. More generally, we say that F is a *slice* if it is a b-slice for some $1 \leq b \leq \frac{n}{2}$.

For example, if X is a finite set, then the function $h : X^7 \to \mathbb{F}$ given by

$$h(x_1, x_2, x_3, x_4, x_5, x_6, x_7) = f(x_2, x_5, x_7) g(x_1, x_3, x_4, x_6)$$

is a slice, and in particular it is a 3-slice with $S = \{2, 5, 7\}$. This leads us to the definition of the multi-slice rank:

Definition 3. Let $X_1, \ldots, X_n$ be finite sets. The *multi-slice rank* of a function

$$F : X_1 \times \cdots \times X_n \to \mathbb{F}$$

is the minimal k such that

$$F = \sum_{i=1}^{k} c_i g_i$$

where the g_i are slices.

It is immediate that any tensor F with slice-rank$(F) = k$ will satisfy multi-slice-rank$(F) \leq k$.

Remark 1. For a finite set X and a field $\mathbb{F}$, consider the function

$$F : X \times X \times X \times X \to \mathbb{F}$$

given by

$$F(x, y, z, w) = \begin{cases} 1 & x = y \text{ and } z = w \\ 0 & \text{otherwise} \end{cases},$$

that is

$$F(x, y, z, w) = \delta(x, y)\delta(z, w)$$

where $\delta(x, y)$ denotes the function that is 1 when $x = y$, and 0 otherwise. Then F will have multi-slice-rank$(F) = 1$ and slice-rank$(F) = |X|$ [12].

Our key observation is that the multi-slice rank of a diagonal tensor is full.

Lemma 2. *Let X be a finite set, and let X^k denote the k-fold Cartesian product of X with itself. Suppose that*

$$F : X^k \to \mathbb{F}$$

is a diagonal tensor, that is

$$F(x_1, \ldots, x_n) = \sum_{a \in A} c_a \delta_a(x_1) \cdots \delta_a(x_n)$$

for some $A \subset X$ where $c_a \neq 0$, and

$$\delta_a(x) = \begin{cases} 1 & x = a \\ 0 & \textit{otherwise} \end{cases}.$$

Then

$$\text{multi-slice-rank}(F) = |A|.$$

Proof. It's evident that the multi-slice rank is at most $|A|$, and so our goal is to prove the lower bound. The proof proceeds by induction on the number of variables. When $n = 2$, this is the usual notion of rank, and so the result follows. Suppose that we can write

$$F(x_1, \ldots, x_n) = \sum_{i=1}^{k} c_i f_i(\overrightarrow{x_{S_i}}) g_i\left(\overrightarrow{x}_{\{1,\ldots,n\}\setminus S_i}\right)$$

where $c_i \neq 0$ and $k < |A|$, and suppose without loss of generality that $|S_i| \leq \frac{n}{2}$ for each i. If there is no i such that $|S_i| = 1$, then choose an arbitrary variable, say x_1, and average over than coordinate. Then we have that

$$\sum_{x_1 \in X} F(x_1, \ldots, x_n) = \sum_{a \in A} c_a \delta_a(x_2) \cdots \delta_a(x_n) = \sum_{i=1}^{k} c_i \tilde{f}_i(\overrightarrow{x_{S_i \setminus \{1\}}}) \tilde{g}_i\left(\overrightarrow{x}_{\{2,\ldots,n\}\setminus S_i}\right),$$

for functions $\tilde{f}, \tilde{g}$ given by averaging f, g over x_1. This contradicts the inductive hypothesis since $\sum_{a \in A} c_a \delta_a(x_2) \cdots \delta_a(x_n)$ will have multi-slice rank equal to $|A| > k$. Now, suppose that there exists some S_i such that $|S_i| = 1$. Then $S_i = \{j\}$ for some $j \in \{1, \ldots, n\}$. Let U be the set of indices for which we have $S_u = \{j\} = S_i$ for each $u \in U$, so that in particular $i \in U$. Define

$$V = \left\{ h : X \to \mathbb{F} : \sum_{x_j \in X} f_u(x_j) h(x_j) = 0 \text{ for all } u \in U \right\}.$$

This vector space has dimension at least $|X| - |U|$. Let $v \in V$ have maximal support, and set $\Sigma = \{x \in X : \ v(x) \neq 0\}$. Then $|\Sigma| \geq \dim V \geq |X| - |U|$, since otherwise there exists nonzero $r \in V$ vanishing on Σ, and the function $v + r$ would have a larger support than v. Multiplying

both sides of our expression by $v(x_j)$ and summing over x_j reduces our tensor dimension by 1. Indeed

$$\sum_{x_j \in X} v(x_j) F(x_1, \dots, x_n) = \sum_{a \in A} c_a \delta_a(x_1) \cdots \delta_a(x_{j-1}) \delta_a(x_{j+1}) \cdots \delta_a(x_n) \left(\sum_{x_j \in X} v(x_j) \delta_a(x_j) \right),$$

and since the sum $\sum_{x_j \in X} v(x_j)\delta_a(x_j)$ will be non-zero for at least $|X| - |U|$ values of $a \in X$, the multi-slice rank of the above must be at least $|A| - |U|$ by the inductive hypothesis. Since

$$\sum_{x_j \in X} v(x_j) f_i(\overrightarrow{x_{S_i}}) = 0$$

for each $i \in U$, it follows that

$$\sum_{x_j \in X} v(x_j) \sum_{i=1}^{k} c_i f_i(\overrightarrow{x_{S_i}}) g_i\left(\overrightarrow{x}_{\{1,\dots,n\} \setminus S_i}\right)$$

will be a sum of at most $k - |U|$ slices, and hence it has multi-slice rank at most $k - |U|$. This implies that $|A| - |U| < k - |U|$, which is a contradiction, and the lemma is proven. □

We conclude this section with the following lemma:

Lemma 3. *Let X be a finite set, and suppose that $F : X^k \to \mathbb{F}$ has tensor rank T and $G : X^k \to \mathbb{F}$ has multi-slice rank R, then*

$$\text{multi-slice-rank}(G \cdot T) \leq RT.$$

Proof. This follows immediately from writing F as a sum of T slices and G as a sum of R rank 1 functions, since the product results in at most RT slices. □

3. The Main Theorems

Let $\mathcal{S}$ denote the set of self-orthogonal vector. Our goal is to find a function

$$I_k : \mathcal{S}^k \to \mathbb{F}_q,$$

of low multi-slice rank satisfying

$$I_k(x_1, \dots, x_k) = \begin{cases} c_1 & x_1, \dots, x_k \text{ are mutually orthogonal and distinct} \\ c_2 & x_1 = \cdots = x_k \\ 0 & \text{otherwise} \end{cases} \tag{3.1}$$

where $c_2 \neq 0$. If $E \subset \mathcal{S}$ is a set without any distinct mutually orthogonal k-tuples, then when we restrict I_k to E^k, I_k will be a diagonal tensor taking the value c_2 on the diagonal, and so by lemma 2, I_k must have multi-slice rank at least $|E|$. It then follows that any set $E \subset \mathcal{S}$ of size

$$E > \text{multi-slice-rank}(I_k),$$

must contain a k-tuple of k distinct mutually orthogonal vectors.

Define

$$F_k : \mathcal{S}^k \to \mathbb{F}_q$$

by

$$F_k(x_1, \dots, x_k) = \prod_{j < l \leq k} \left(1 - \langle x_j, x_l \rangle^{q-1}\right). \tag{3.2}$$

For example,

$$F_3(x, y, z) = \left(1 - \left(\sum_{i=1}^{n} x_i y_i\right)^{q-1}\right) \left(1 - \left(\sum_{i=1}^{n} y_i z_i\right)^{q-1}\right) \left(1 - \left(\sum_{i=1}^{n} z_i x_i\right)^{q-1}\right),$$

is a polynomial of degree $3(q-1)$ in n variables. This function satisfies

$$F(x_1,\dots,x_k)=\begin{cases}1 & x_1,\dots,x_k \text{ are mutually orthogonal}\\ 0 & \text{otherwise}\end{cases}.$$

We note that $x_1,\dots,x_k$ are not required to be unique for $F(x_1,\dots,x_k)$ to equal 1, and in particular if $x_1=\cdots=x_k$, then $F_k(x_1,\dots,x_k)=1$ since $\langle x,x\rangle=0$ for every $x\in\mathcal{S}$. Ideally, we would like to take $I_k=F_k$, however this only works when $k=2$, and does not hold due to the repetitions for $k\geq 3$. For example, when $k=4$, if $x,y\in\mathbb{F}_q^n$ satisfy$\langle x,y\rangle=0$, then

$$F_4(x,x,y,y)=1,$$

and we need to try to modify F_k to make it zero in this case.

3.1. **Constructing an indicator function for distinct k-tuples.** Let X be a finite set, $\mathbb{F}$ a field, and let $X^k=X\times\cdots\times X$ denote the Cartesian product of X with itself k times. For every $\sigma\in S_k$, define

$$f_\sigma: X\times\cdots\times X\to\mathbb{F}$$

to be the function that is 1 if $(x_1,\dots,x_k)$ is a fixed point of σ, and 0 otherwise.

Lemma 4. *We have the identity*

$$\sum_{\sigma\in S_k}\operatorname{sgn}(\sigma)f_\sigma(x_1,\dots,x_k)=\begin{cases}1 & \textit{if } x_1,\dots,x_k \textit{ are distinct}\\ 0 & \textit{otherwise}\end{cases},$$

where $\operatorname{sgn}(\sigma)$ *is the sign of the permutation.*

Proof. By definition,

$$\sum_{\sigma\in S_k}\operatorname{sgn}(\sigma)f_\sigma(x_1,\dots,x_k)=\sum_{\sigma\in\operatorname{Stab}(\overrightarrow{x})}\operatorname{sgn}(\sigma)$$

where $\operatorname{Stab}(\overrightarrow{x})\subset S_k$ is the stabilizer of $\overrightarrow{x}$. Since the stabilizer is a product of symmetric groups, this will be non-zero precisely when $\operatorname{Stab}(\overrightarrow{x})$ is trivial, and hence $x_1,\dots,x_k$ must be distinct. This vector is then fixed only by the identity element, and so the sum equals 1. □

Lemma 5. *Let $\mathcal{C}_1\subset S_k$, $\mathcal{C}_1=Cl(1\ 2\ \cdots\ k)$, be the conjugacy class of the k-cycles in S_k, and define*

$$H_k(x_1,\dots,x_k)=\sum_{\substack{\sigma\in S_k\\ \sigma\notin\mathcal{C}_1}}\operatorname{sgn}(\sigma)f_\sigma(x_1,\dots,x_k).$$

Then

$$H_k(x_1,\dots,x_k)=\begin{cases}1 & \textit{if } x_1,\dots,x_k \textit{ are distinct}\\ (-1)^k(k-1)! & \textit{if } x_1=\cdots=x_k\\ 0 & \textit{otherwise}\end{cases}.$$

Proof. This follows from lemma 4 and the fact that the conjugacy class of the k-cycle in S_k has $(k-1)!$ elements each with sign equal to $(-1)^{k-1}$. □

This function can be used to zero-out those tuples of vectors with repetitions, and so when $X=\mathbb{F}_q^n$, we define

$$I_k(x_1,\dots,x_k):=F_k(x_1,\dots,x_k)H_k(x_1,\dots,x_k),$$

and this will satisfy precisely the desired properties of equation (3.1) with $c_1=1$ and $c_2=(-1)^k(k-1)!$. Since we want $c_2\neq 0$, this imposes the constraint $p\geq k$.

When $k=2$,

$$H_2(x_1,x_2)=1,$$

when $k=3$

$$H_3(x_1,x_2,x_3)=1-\delta(x_1,x_2)-\delta(x_2,x_3)-\delta(x_3,x_2)$$

and when $k = 4$

$$\begin{aligned} H_4(x_1, x_2, x_3, x_4) =& 1 - \delta(x_1, x_2) - \delta_2(x_2, x_3) - \delta(x_3, x_4) - \delta(x_4, x_1) - \delta(x_1, x_3) - \delta(x_2, x_4) \\ &+ 2\delta(x_1, x_2, x_3) + 2\delta(x_2, x_3, x_4) + 2\delta(x_3, x_4, x_1) + 2\delta(x_4, x_1, x_2) \\ &+ \delta(x_1, x_2)\delta(x_3, x_4) + \delta(x_1, x_3)\delta(x_2, x_4) + \delta(x_1, x_4)\delta(x_2, x_3) \end{aligned}$$

where

$$\delta(x_1, \dots, x_k) = \begin{cases} 1 & x_1 = \cdots = x_k \\ 0 & \text{otherwise} \end{cases}.$$

As a function on X^3, slice-rank$(\delta(x_1, x_2)) = 1$, but on X^4 the function

$$\delta(x_1, x_2)\delta(x_3, x_4)$$

has slice rank equal to $|X|$. This function *does* however have multi-slice rank equal to 1. When writing $H_k(x_1, \dots, x_k)$ as a linear combination of δ functions, it will always have multi-slice rank at most $k!$, since it is a sum of $k!$ slices. However, starting from $k = 4$ there will be functions in the sum whose slice rank is maximal, and this is why the multi-slice rank is needed to prove theorems 1 and 2 for $k \geq 4$ and $k \geq 3$ respectively.

3.2. **Proof of Theorem 1.**

Proof. Let F_k be defined as in (3.2), and let $H_k : \mathcal{S}^k \to \mathbb{F}_q$ be defined as in lemma 5. Define

$$I_k(x_1, \dots, x_k) = F_k(x_1, \dots, x_k)H_k(x_1, \dots, x_k).$$

This function satisfies

$$I_k(x_1, \dots, x_k) = \begin{cases} 1 & x_1, \dots, x_k \text{ are mutually orthogonal and distinct} \\ (-1)^k(k-1)! & \text{when } x_1 = x_2 = \cdots = x_k \\ 0 & \text{otherwise} \end{cases},$$

and so we assume that $q = p^r$ where $p \geq k$, since otherwise the above function would be zero when $x_1 = \cdots = x_k$. It then follows from lemma 2 that any set $E \subset \mathcal{S}$ satisfying

$$|E| > \text{multi-slice-rank}(I_k)$$

must contain a k-tuple of distinct mutually orthogonal vectors. In what follows, we bound the multi-slice rank of I_k. For any t, we may write

$$F_t(x_1, \dots, x_j) = \prod_{j < l \leq t} \left(1 - \left(\sum_{i=1}^{n} x_{ji}x_{li}\right)^{q-1}\right),$$

and expand this product as a polynomial. The total number of terms in the expansion of this polynomial will be at most the total number of terms in the polynomial

$$(1 + z_1 + \cdots + z_n)^{(q-1)\frac{t(t-1)}{2}}.$$

Letting $\beta = \frac{t(t-1)}{2}(q-1)$, it is a classical combinatorics result that the above has at most

$$\binom{n+\beta}{\beta} \leq \frac{1}{\beta!}(n+\beta)^{\beta}$$

terms. Since this was a bound on the total number of terms, this bounds the tensor rank of F_t, and by lemma 3 combined with the fact that multi-slice-rank$(H_k) \leq k!$, we immediately retrieve the bound

$$\text{multi-slice-rank}(I_k) \leq k!\binom{n+\beta}{\beta}.$$

By being slightly more careful, we can obtain significantly better results. For $\sigma \in S_k$ let f_σ be defined as in subsection 3.1, and consider

$$f_\sigma(x_1, \dots, x_k)F_k(x_1, \dots, x_k).$$

Let t be the number of cycles in the permutation decomposition of σ, and let r denote the number of fixed points. Then f_σ will be a product of $t-r$ multi-variable delta functions. Suppose that $f_\sigma(x_1,\dots,x_k)=1$. Then there will be exactly t distinct elements among the x_i, and so we have the equality

$$f_\sigma(x_1,\dots,x_k)F_k(x_1,\dots,x_k)=f_\sigma(x_1,\dots,x_k)F_t(\overrightarrow{x}_S)$$

where $S\subset\{1,2,\dots,k\}$ is a subset of indices such that no two elements of S are in the same cycle of the permutation σ. By slicing in between the cycles, it follows that the multi-slice rank of this tensor will be at most the slice-rank of the t-dimensional tensor $F_t(z_1,\dots,z_t)$. Thus

$$\text{multi-slice-rank}(f_\sigma(x_1,\dots,x_k)F_t(\overrightarrow{x}_S))\leq\text{multi-slice-rank}(F_t),$$

where the right hand side is viewed as a t-dimensional tensor. This implies the bound

$$\text{multi-slice-rank}(I_k)\leq\sum_{i=2}^{k}\begin{bmatrix}k\\i\end{bmatrix}\text{multi-slice-rank}(F_i) \tag{3.3}$$

where $\begin{bmatrix}k\\i\end{bmatrix}$ denotes the unsigned Stirling numbers of the first kind, which count the number of permutations of length k made of exactly i disjoint cycles, and multi-slice-rank(F_i) refers to the multi-slice rank of F_i when viewed as an i-dimensional tensor. Returning to (3.2), the product form for F_t,

$$F_t(z_1,\dots,z_t)=\prod_{j<l\leq t}\left(1-\left(\sum_{i=1}^{n}z_{ji}z_{li}\right)^{q-1}\right),$$

we may bound the multi-slice rank by counting the number of polynomials in z_1 that can occur. Expanding the above product, it will be a linear combination of terms of the form

$$(z_{11}^{e_{11}}\cdots z_{t1}^{e_{t1}})\cdots(z_{1n}^{e_{1n}}\cdots z_{tn}^{e_{tn}})$$

where $e_{ij}\leq q-1$ for all i,j and

$$\sum_{j=1}^{n}e_{ij}\leq(q-1)(t-1)$$

for every i. It then follows that

$$\text{multi-slice-rank}\,(F_t)\leq\#\left\{\vec{v}\in\{0,1,\dots,q-1\}^n:\ \sum_{i=1}^{n}v_i\leq(q-1)(t-1)\right\}.$$

Let

$$C_t=\#\left\{\vec{v}\in\{0,1,\dots,q-1\}^n:\ \sum_{i=1}^{n}v_i\leq(q-1)(t-1)\right\}.$$

This can be thought of as putting $(q-1)(t-1)$ balls into $n+1$ buckets, where the $(n+1)^{st}$ bucket represents the constant term. By the balls and buckets theorem,

$$C_t\leq\binom{n+(q-1)(t-1)}{(q-1)(t-1)},$$

however this significantly over counts the number of such vectors. Each vector v satisfying

$$(q-1)(t-2)<\sum_{i=1}^{n}v_i\leq(q-1)(t-1)$$

will be counted exactly once. Each vector v satisfying

$$(q-1)(t-3)<\sum_{i=1}^{n}v_i\leq(q-1)(t-2)$$

will be counted exactly $n+1 = \binom{n+1}{1}$ times, since we may put an extra $(q-1)$ balls into any of the $n+1$ buckets and arrive at the same vector. Each vector v satisfying

$$(q-1)(t-r-2) < \sum_{i=1}^{n} v_i \leq (q-1)(t-r-1)$$

will be counted exactly $\binom{n+r}{r}$ times, since we have r extra $(q-1)$'s to put into $n+1$ buckets. Thus

$$\binom{n+(q-1)(k-1)}{(q-1)(k-1)} = C_k + \left(\binom{n+1}{1} - \binom{n+0}{0}\right) C_{k-1} + \left(\binom{n+2}{2} - \binom{n+1}{1}\right) C_{k-2} + \cdots,$$

that is

$$\binom{n+(q-1)(k-1)}{(q-1)(k-1)} = C_k + nC_{k-1} + \binom{n+1}{2} C_{k-2} + \binom{n+2}{3} C_{k-3} + \cdots \tag{3.4}$$

$$= \sum_{j=0}^{k} \binom{n-1+j}{j} C_{k-j}. \tag{3.5}$$

We remark that from this identity, one could use generating series to work out C_k exactly as a sum of binomial coefficients, and hence get an exact expression for

$$\sum_{i=t}^{k} \begin{bmatrix} k \\ t \end{bmatrix} C_t.$$

However this expression has an unpleasant form, and so to have the theorem in a simpler form, we proceed by proving the following lemma:

Lemma 6. *Let* $\begin{bmatrix} k \\ i \end{bmatrix}$ *denote the unsigned Stirling numbers of the first kind and let* $n = \binom{k}{2}$. *Then*

$$\begin{bmatrix} k \\ k-i \end{bmatrix} \leq \binom{n}{i}$$

for all $1 \leq i \leq k$.

Proof. Let $\sigma_1, \ldots, \sigma_n$ denote the $\binom{k}{2}$ transpositions in S_k. Every element of S_k with $k-i$ cycles is the product of exactly i distinct transpositions. Hence

$$\begin{bmatrix} k \\ k-i \end{bmatrix} \leq \binom{n}{i},$$

and the lemma is proven. □

By equation (3.4) and the above lemma, for $n \geq \binom{k}{2}$ it follows that

$$\binom{n+(q-1)(k-1)}{(q-1)(k-1)} \geq \sum_{t=2}^{k} \begin{bmatrix} k \\ t \end{bmatrix} C_t,$$

and so

$$\text{multi-slice-rank}(I_k) \leq \binom{n+(q-1)(k-1)}{(q-1)(k-1)},$$

and the proof is complete. □

3.3. **Right Angles in $\mathbb{F}_q^n$.** In this subsection, we prove theorem 2.

Proof. Define

$$J_k(x_1,\dots,x_k,x_{k+1}) = H_{k+1}(x_1,\dots,x_{k+1})F_k(x_1-x_{k+1},x_2-x_{k+1},\dots,x_k-x_{k+1})$$

where H_{k+1} and F_k are defined in lemma 5 and equation (3.2), respectively. Then

$$J_k(x_1,\dots,x_k,x_{k+1}) = \begin{cases} 1 & \text{if } x_1,\dots,x_{k+1} \text{ are distinct and form a right corner} \\ (-1)^{k+1}k! & \text{if } x_1=\cdots=x_{k+1} \\ 0 & \text{otherwise} \end{cases}.$$

If $A \subset \mathbb{F}_q^n$ contains no k-right corners, then the restriction of J_k to A^{k+1} will be a diagonal tensor with $(-1)^{k+1}k!$ on the diagonal, and so lemma 2 implies that

$$|A| \leq \text{multi-slice-rank}(J_k)$$

as long as $q = p^r$ where $p > k$ so that $k!$ is not equal to 0. All that remains is to bound the multi-slice rank of J_k. For $\sigma \in S_{k+1}$ let f_σ be defined as in subsection 3.1, and consider

$$f_\sigma(x_1,\dots,x_k)F_k(x_1-x_{k+1},x_2-x_{k+1},\dots,x_k-x_{k+1}).$$

Let t be the number of cycles in the permutation decomposition of σ and suppose that $f_\sigma(x_1,\dots,x_{k+1}) = 1$. Then there will be exactly t distinct elements among the x_i, and so $F_k(x_1-x_{k+1},x_2-x_{k+1},\dots,x_k-x_{k+1})$ can be rewritten as a function of t variables, and we have the equality

$$f_\sigma(x_1,\dots,x_{k+1})F_k(x_1-x_{k+1},x_2-x_{k+1},\dots,x_k-x_{k+1}) = f_\sigma(x_1,\dots,x_{k+1})F_{t-1}(\overrightarrow{x}_S - \overrightarrow{x}_{k+1})$$

where $S \subset \{1,2,\dots,k\}$ is a subset of $t-1$ indices. In particular, whichever cycle $k+1$ falls into, we take it to be the representative variable of that cycle. For convenience of notation, let

$$G_t(z_1,\dots,z_t) = F_{t-1}(z_1-z_t,\dots,z_{t-1}-z_t).$$

By slicing in between the different cycles, it follows that the multi-slice rank of this tensor will be at most the slice-rank of the t-dimensional tensor $G_t(z_1,\dots,z_t)$. By equation (3.2),

$$G_t(z_1,\dots,z_t) = \prod_{j<l\leq t-1}\left(1-\left(\sum_{i=1}^n (z_{ji}-z_{ti})(z_{li}-z_{ti})\right)^{q-1}\right),$$

and we may expand this as a polynomial and count the number of polynomials in z_1 that can occur. As z_1 appears in at most $(q-1)(t-2)$ terms in the product, and each term has degree at most 2, we have that

$$\begin{aligned}\text{multi-slice-rank}(G_t(z_1,\dots,z_t)) &\leq \#\left\{\vec{v}\in\{0,1,\dots,q-1\}^n : \sum_{i=1}^n v_i \leq (q-1)(t-2)\right\} \\ &= C_{t-1}.\end{aligned}$$

Thus

$$\text{multi-slice-rank}(J_k) \leq \sum_{t=1}^{k+1}\begin{bmatrix} k+1 \\ k+1-t\end{bmatrix} C_{t-1}.$$

By lemma 6, for $n \geq \binom{k+1}{2}$ we have that

$$\text{multi-slice-rank}(J_k) \leq \sum_{t=1}^{k+1}\binom{n}{t}C_{t-1} \leq \sum_{j=0}^{k}\binom{n-1+j}{j}C_{k-j},$$

and so by equation (3.4) it follows that

$$\text{multi-slice-rank}(J_k) \leq \binom{n+(q-1)(k-1)}{(q-1)(k-1)},$$

and the theorem is proven. □

Acknowledgments

I would like to thank Will Sawin for his helpful comments and for his simple proofs of lemmas 4 and 5. This work was partially supported by the NSERC PGS-D scholarship, and by Ben Green's ERC Starting Grant 279438, Approximate Algebraic Structure and Applications.

References

1. Michael Bennett, *Right Angles in $\mathbb{F}_q^d$*, preprint, 2015, arXiv:1511.08942.
2. Jonah Blasiak, Thomas Church, Henry Cohn, Joshua Grochow, Eric Naslund, William F. Sawin, and Christopher Umans, *On cap sets and the group-theoretic approach to matrix multiplication*, preprint, 2016, arXiv:1605.06702.
3. Ernie Croot, Vsevolod Lev, and Peter Pach, *Progression-free sets in $\mathbb{Z}_4^n$ are exponentially small*, Ann. of Math. (2) **185** (2017), 331–337, arXiv:1605.01506, doi:10.4007/annals.2017.185.1.7.
4. Jordan Ellenberg and Dion Gijswijt, *On large subsets of $\mathbb{F}_q^n$ with no three-term arithmetic progression*, Ann. of Math. (2) **185** (2017), 339–343, arXiv:1605.09223, doi:10.4007/annals.2017.185.1.8.
5. Paul Erdős, Abraham Ginzburg, and Abraham Ziv, *A theorem in additive number theory*, Bull. Res. Council Israel **10F** (1961), 41–43.
6. Paul Erdős and Endre Szemerédi, *Combinatorial properties of systems of sets*, J. Combinatorial Theory Ser. A **24** (1978), 308–313, doi:10.1016/0097-3165(78)90060-2.
7. Gennian Ge and Chong Shuangguan, *Rank counting and maximum subsets of $\mathbb{F}_q^n$ containing no right angles*, preprint, 2016, arXiv:1612.08255.
8. Alex Iosevich and Steven Senger, *Orthogonal systems in vector spaces over finite fields*, Electron. J. Combin. **15** (2008), no. 1, Research Paper 151, 10. MR 2465775, arXiv:0807.0592.
9. Eric Naslund, *Exponential bounds for the Erdős-Ginzburg-Ziv constant*, preprint, 2017.
10. Eric Naslund and William F. Sawin, *Upper bounds for sunflower-free sets*, preprint, 2016, arXiv:1606.09575.
11. Terence Tao, A symmetric formulation of the Croot–Lev–Pach–Ellenberg–Gijswijt capset bound, blog post, 2016, `http://terrytao.wordpress.com/2016/05/18/a`.
12. Terence Tao and William F. Sawin, Notes on the "slice rank" of tensors, blog post, 2016, `https://terrytao.wordpress.com/2016/08/24/notes-on-the-slice-rank-of-tensors/`.
13. Le Anh Vinh, *On the number of orthogonal systems in vector spaces over finite fields*, Electron. J. Combin. **15** (2008), no. 1, Note 32, 4. MR 2438588, arXiv:0807.2690.

E-mail address: `naslund@princeton.edu`

Fine Hall, Washington Road; Princeton NJ 08544